

SA 21.5: A 33GB/s 13.4Mb Integrated Graphics Accelerator and Frame Buffer

R. Torrance, I. Mes, B. Hold, D. Jones, J. Crepeau,
P. DeMone, D. MacDonald, C. O'Connell, P. Gillingham,
R. White¹, S. Duggins², D. Fielder²

MOSAID Technologies Inc., Kanata, Ontario, Canada

¹Accelerix Incorporated, Ottawa, Ontario, Canada

²Symbionics Ltd., Cambridge, England, UK

With the increasing capacity of commodity DRAM, a dichotomy has appeared between memory requirements of PC graphics systems, and standard DRAM for PC main memory. PC graphics subsystems require wide, relatively shallow memory for high data bandwidth while commodity DRAM is too narrow and too deep. Integrating DRAM frame buffer and graphics accelerator logic on the same chip solves this problem (Figure 1).

Reported integrated DRAM and logic devices separate DRAM and ASIC logic portions with a traditional memory interface [1]. Although this approach has a number of benefits over the discrete solution, much greater improvements are available by more tightly integrating the DRAM and logic, something not possible at the board level. This device integrates parts of the graphics processor within the DRAM to increase performance. Figure 2 shows the architecture of one bank of the frame buffer, where the pixel processing unit (PPU) and the serial output registers (SORs) are integrated into the DRAM architecture [2]. This allows the bus width between the DRAM frame buffer and the processor to be 4096b.

Figure 3a shows a traditional DRAM, where a relatively narrow databus runs the length of the sense amp straps. To allow for massively-parallel DRAM access, the frame buffer uses databuses orthogonal to the standard direction (Figure 3b). These databuses can be routed in metal2 (M2) over an array with capacitor-over-bitline processing. This allows up to 1 databus for each column of DRAM, depending on the M2 pitch of the process. The device employs 8 banks, each of 512b databus width.

The PPU is 512b wide per bank, requiring one bit of the PPU to pitch-match to 4 DRAM columns. With this architecture, the choice of circuits included in this massively parallel processor must be judicious. It is limited to the most basic pixel operations: the raster operations. Figure 4 shows a single bit of the PPU, and how it is pitch-matched to the DRAM to accelerate raster operations. The PPU registers are built using dual port 6T SRAM cells as shown in Figure 5a, while the rastop function unit is implemented with a pMOS based 8-1 multiplexor as shown in Figure 5b. This allows the PPU to perform any of the 256 rastops possible with 3 input variables. All 512b are identical, and since there is no dataflow directly between bits, redundant PPU bits can be easily added (pitch-matched to the DRAM redundant columns). This allows more aggressive core rules for the PPU to reduce area.

This frame buffer architecture allows the graphics controller to accelerate some of the most common graphic operations. For instance, a block move can require a source pixel read, a destination pixel read, a rastop, and a destination pixel write. In a typical system with a 64b memory interface, this is done 64b at a time. In this chip, all reads can be done at up to 4096b wide, as can the actual raster operation, and the write back. For operations that require realignment of source pixels to destination

pixels, a 32b funnel shifter is built under the FB data-bus in each bank, allowing realignment of 256b at a time. Also, a 32b word can be written to all 16 PPU words simultaneously. If all 8 banks are enabled, this allows for a 4096b write to DRAM in a single cycle, useful for screen clearing.

The gate array logic incorporates the graphic accelerator, PCI interface, VGA core and video input block to provide the necessary system level functionality for a PC graphics device. Within the accelerator logic the highly-parallel nature of the DRAM and its embedded logic gives a much larger control space than a normal external D/VRAM interface. This requires a novel approach to BITBit operation since data alignment, manipulation (ROP) and masking must be performed in the most optimal way to achieve the best performance. The availability of 5 512B registers within the PPU allows data to be cached for some operations to increase performance further. The custom-designed pixel output path (POP) logic includes a 64b hardware cursor, a video scaler with a 4-tap FIR filter for horizontal scaling, and 2-tap for vertical scaling, color space conversion circuits, and a 135MHz LUTDAC. This device uses dual fully integrated PLLs, one for the 66MHz main clock (MCLK), and one for the 135MHz pixel clock (PCLK).

Noise is a concern on a mixed DRAM/logic chip. Noise from the ASIC logic can cause a degradation of the refresh time for DRAM cells. Refresh time is programmable, to allow the use of the maximum refresh time possible after processing. The device employs separate power pins and on-chip power rails for DRAM, gate array, PLLs, DACs, and POP logic. Critical circuitry is guarded with substrate and well guard rings.

Testing methodologies have yet to be standardized for mixed DRAM/logic chips. Although BIST is useful for embedded SRAMs, redundancy and complicated cell coupling tests make it difficult to use for large embedded DRAMs. BIST also complicates the debug of an embedded DRAM, since the controllability and observability of that DRAM is severely limited. In this device all DRAM data and control is multiplexed onto the device pins in test mode, using an on-chip test interface. This circuit makes the entire frame buffer appear similar to a standard asynchronous DRAM in memory test, allowing use of standard DRAM test software. All bits of the SORs and PPUs are mapped into unused row addresses for test. Even the rastop processor is mapped to unused row addresses to allow the entire frame buffer to be easily tested on a standard DRAM tester. The digital logic blocks are tested using full scan methodology and functional test vectors, and the test interface allows access for these paths. For device burn-in only, the logic incorporates a BIST block for the DRAM array. This ensures that the DRAM is cycled during post-package burn-in testing with only the requirement for power, ground, and a few control pins connected to a tester.

The device is implemented in a 0.35µm/0.5µm blended DRAM and logic process using triple-level-metal. Table 1 shows the features of the chip. Figure 6 is a chip micrograph.

Acknowledgments:

The authors acknowledge the support of M. Liang and C. Yoo of TSMC in developing this chip.

References:

- [1] Miyano, S., et al., "A 1.6GB/s Data Transfer Rate 8Mb Embedded DRAM," IEEE Journal of Solid-State Circuits, vol. 30, no. 11, pp. 1281-1285, Nov., 1995.
- [2] Elliott, D., et al., "Computational RAM: A Memory SIMD Hybrid and Its Application to DSP," Custom Integrated Circuits Conference, May, 1992.

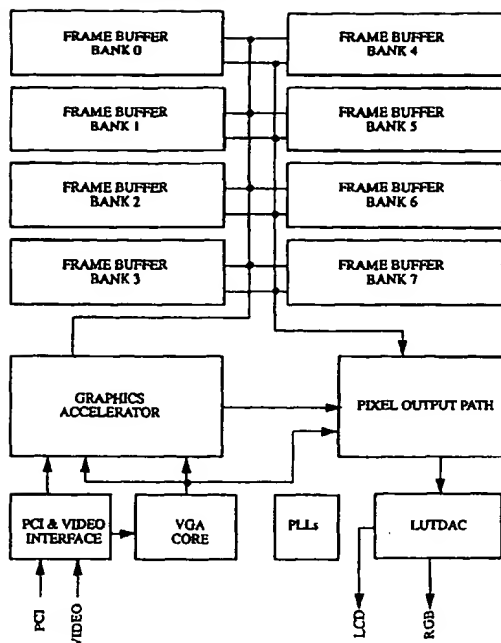


Figure 1: Chip architecture.




Figure 2: Frame buffer bank. Figure 3a: Traditional DRAM.




Figure 3a: Traditional DRAM.




Figure 3b: Wide databus scheme. Figure 4: See page 461.




Figure 5a: PPU register bit.




Figure 5b: RASTOP processor. Figure 6 and Table 1: See page 461.

Maximum screen sizes	1280x1024x8b/pixel 1024x768x16b/pixel 800x600x24b/pixel
Technology	0.35µm/0.5µm BlendIC, 3M
DRAM	13.4Mb
Pitch-matched logic	160k gates
SRAM	38kb
Maximum system clock rate	66MHz
Maximum pixel clock rate	135MHz
Supply voltage	3.3V
Package	208-pin QFP

Table 1: Chip overview.




Figure 6: Chip micrograph.




Figure 4: Pixel processing unit.